

Finite mixture model of conditional dependencies modes to cluster categorical data.

Matthieu Marbac^a, Christophe Biernacki^b, Vincent Vandewalle^c

^a*DGA & Inria Lille & University Lille 1*

^b*University Lille 1 & CNRS & Inria Lille*

^c*EA 2694 University Lille 2 & Inria Lille*

Abstract

We propose a parsimonious extension of the classical latent class model to cluster categorical data by relaxing the class conditional independence assumption. Under this new mixture model, named Conditional Modes Model, variables are grouped into conditionally independent blocks. The corresponding block distribution is a parsimonious multinomial distribution where the few free parameters correspond to the most likely modality crossings, while the remaining probability mass is uniformly spread over the other modality crossings. Thus, the proposed model allows to bring out the intra-class dependency between variables and to summarize each class by a few characteristic modality crossings. The model selection is performed via a Metropolis-within-Gibbs sampler to overcome the computational intractability of the block structure search. As this approach involves the computation of the integrated complete-data likelihood, we propose a new method (exact for the continuous parameters and approximated for the discrete ones) which avoids the biases of the BIC criterion pointed out by our experiments. Finally, the parameters are only estimated for the best model via an EM algorithm. The characteristics of the new model are illustrated on simulated data and on two biological data sets. These results strengthen the idea that this simple model allows to reduce biases involved by the conditional independence assumption and gives meaningful parameters. Both applications were performed with the R package **CoModes**¹ where the proposed model is implemented.

Keywords: categorical data, clustering, Metropolis-within-Gibbs sampler,

¹Downloadable at https://r-forge.r-project.org/R/?group_id=1809

1. Introduction

Clustering [1] is an important tool for practitioners confronted with a complex data set. Indeed, this method allows to extract main information from data by grouping individuals into *homogeneous* classes. This paper focuses on categorical variables clustering. Even if these variables are little informative, they are present in many different fields (biology, sociology, marketing...) because they are usually easily accessible.

Clustering methods can be split into two approaches: the *geometrical* ones based on the distances between individuals and the *probabilistic* ones which model the data generation.

If these first approaches are generally simpler and faster than the others, they are sensitive to the used distance between individuals. Geometrical approaches, to cluster categorical data, either define a metric in the initial variable space like the *k-means* [2], either compute their metric on the axes of the multiple correspondence analysis [3, 4]. Indeed, these approaches consider that the classes are *homogeneous* when the distance between the individuals of the class and its center is small. However, lots of geometrical approaches can be interpreted as probabilistic ones [5] revealing probabilistic hidden assumptions made by the geometrical ones. Moreover, the probabilistic approaches allow to solve difficult questions, like the class number selection, in a rigorous mathematical framework.

The probabilistic approaches consider that the classes are *homogeneous* when the individuals of a class are drawn from the same distribution. Thus, finite mixture models, which are the most classical of the probabilistic approaches, meet this objective by approaching the data distribution with a finite mixture of parametric distributions [6]. In addition, the obtained partition is meaningful since each class is described by the parameters of the corresponding component.

The most widely used mixture model to cluster categorical data sets is the latent class model [7, 8, 9], which assumes the conditional independence between variables. In this article, we refer to this model as *Conditional Independence Model* (further mentioned as CIM). In this model, as the classes

are explicitly described by the probability of each modality for each variable, the interpretation is easy. Moreover, the sparsity caused by the conditional independence assumption is a great advantage since it circumvents the curse of dimensionality. In practice, this model obtains good results in lots of applications [10, 11, 12]. However, applications [13] can show that CIM overestimates the class number when the conditional independence assumption is violated (see also our experiments presented in Section 5). Furthermore, the larger is the number of variables, the higher is the risk to observe conditionally correlated variables in a data set, and consequently the higher is the risk to involve such biases by using CIM.

Different models relax the conditional independence assumption. Among them, the *multilevel latent class model* [14, 15] assumes that the conditional dependency between the observed variables can be explained by other unobserved variables. This model has connections with the approach modeling the intra-class dependencies by using a latent continuous variable and a probit function [16]. Recently proposed, the *mixture of latent trait analyzers* [17, 18] is a good challenger for CIM since it assumes that the categorical variable distribution depends on many latent variables: one categorical variable (the class) and many continuous latent variables (modeling the intra-class dependencies between the observed categorical variables). However, the parameters are hardly estimated directly, so the authors use a variational approach. Furthermore, the intra-class dependencies can be hardly interpretable by the practitioner, since the correlations are interpreted according to relationships with unobserved continuous variables.

The log-linear models [19, 20] purpose is to model the individual log-probability by selecting interactions between variables. Thus, the most general mixture model is the *log-linear mixture model* where all the kinds of interactions can be considered. It has been used for a long time [21] and it obtains good results in many applications like the clustering of radiographic cross-diagnostics [22] or in a market segmentation [13]. However this model family is huge and the model selection is a real challenge. In the literature, authors fix by advance the modeled interactions or they perform a deterministic search like the *forward* method which is sub-optimal. Furthermore, the number of parameters increases with the conditional modality crossings, so there is an over-fitting risk and the interpretation becomes harder.

In this paper, we propose a sparse mixture model relaxing the conditional independence assumption to overcome the biases caused by CIM. This new

model, named *Conditional Modes Model* (referred in this article by CMM), groups the variables into *conditionally independent blocks*, allowing to consider the main conditional dependencies. Such an idea was already proposed to cluster continuous and categorical data in the Multimix software [23, 24]. However, the specific distribution of the block that we adopt here is a multinomial per modes distribution which assumes that few modality crossings, named *modes*, are characteristic and that the other ones follow a uniform distribution. Thus, the associated multinomial distribution is parsimonious, its free parameters being limited to the few parameters of the modes.

This simple mixture model (CMM) is a good challenger for the mixture model with conditional independence assumption (CIM), since it preserves the sparsity and avoids many biases through modeling of the main conditional correlations. It can be also interpreted as a parsimonious version of the log-linear mixture model. Indeed, the repartition of the variables into blocks defines the considered interactions while the mode distribution into blocks defines a specific distribution for each interaction. Furthermore, resulting classes are meaningful since the intra-class dependencies are brought out at two complementary levels: the block variable interaction level and the associated mode interaction level (through locations and probabilities). Note that CMM is a comprehensive approach since it includes CIM and a part of its parsimonious versions [8].

For a fixed model (class number, repartition of the variables into blocks and mode numbers), the maximum likelihood estimate is obtained via an EM algorithm. The model selection is performed via a Metropolis-within-Gibbs sampler generating a new block variable repartition into blocks and new mode numbers by a Metropolis-Hastings step. It is performed for a fixed number of classes and avoids combinatorial problems involved by the selection of the blocks of variables and of the mode number. It is based on the fact that the integrated complete-data likelihood, required for the acceptance probability computation of the Metropolis-Hastings inside the Gibbs sampler, is accessible and non ambiguous through weakly informative conjugate prior. Finally, this approach has two main advantages. It allows to reduce the bias of the BIC-like approach (the overestimation of the number of modes by this approach is illustrated during our numerical experiments). Furthermore, it allows to perform an efficient model selection in a reasonable computational time since the parameters are only estimated for the unique selected model. Thus, this approach is a possible answer to the combinatorial model selection problem which is known to be a real challenge for a log-linear

mixture model.

This paper is organized as follows. Section 2 presents the Conditional Modes Model. Section 3 is devoted to maximum likelihood estimation via an EM algorithm. Section 4 presents the Metropolis-within-Gibbs sampler performing the model selection through the integrated complete-data likelihood. In Section 5, we show that the proposed approach computing the integrated complete-data likelihood sharply reduces the biases of the BIC-like approach and we numerically underline both the good behavior of the Metropolis-within-Gibbs sampler and the flexibility of CMM on simulated data. Section 6 presents two clusterings of biological data sets performed by the R package **CoModes**². A conclusion is drawn and future extensions are discussed in Section 7.

2. Conditional modes model

2.1. Conditional modes model framework

Observations are described with B categorical variables $\mathbf{x} = (\mathbf{x}^1, \dots, \mathbf{x}^B)$ using the complete disjunctive coding, where \mathbf{x}^b has m_b modalities. Let a partition $\boldsymbol{\sigma} = (\boldsymbol{\sigma}_1, \dots, \boldsymbol{\sigma}_d)$ of $\{1, \dots, B\}$ determining a repartition of variables in d blocks. The j -th block is also denoted by

$$\mathbf{x}^j = \{\mathbf{x}^b; b \in \boldsymbol{\sigma}_j\}. \quad (1)$$

We still adopt the disjunctive coding $x^{jh} = 1$ if the individual takes the modality crossing h and $x^{jh} = 0$ otherwise, where $h \in \{1, \dots, m_j\}$. Thus, \mathbf{x}^j corresponds to a new categorical variable having m_j modalities, m_j being the number of modality crossings of the initial variables affected into the block j determined by $m_j = \prod_{b \in \boldsymbol{\sigma}_j} m_b$.

The proposed model, below referred as Conditional Modes Model (CMM), assumes that data arise independently from a mixture of g components of *independent blocks of initial variables*, the intra-block distribution being *multinomial per modes*. In each class, the multinomial distribution of \mathbf{x}^j is also assumed to have few free parameters in comparison to the number of modalities m_j and corresponding to the *modes* of the distribution. More precisely, they are defined as the locations of the largest probabilities, while the other

²Downloadable at https://r-forge.r-project.org/R/?group_id=1809

parameters are equal. A particular model is denoted by $\omega = (g, \sigma, \ell)$ where $\ell = (\ell_1, \dots, \ell_g)$ groups all the mode numbers with $\ell_k = (\ell_{k1}, \dots, \ell_{kd})$, ℓ_{kj} being the number of modes of \mathbf{x}^j for the class k ($0 < \ell_{kj} < m_j$).

2.2. Probability distribution functions of the conditional modes model

Using $p(\cdot; \cdot)$ as a generic notation for the probability distribution function (pdf), for a known model ω , the CMM's pdf can be written as

$$p(\mathbf{x}; \theta, \omega) = \sum_{k=1}^g \pi_k p(\mathbf{x}; \alpha_k, \sigma, \ell_k), \quad (2)$$

where $\theta = (\pi, \alpha)$ denotes the whole mixture parameters: $\pi = (\pi_1, \dots, \pi_g)$ is the vector of class proportions with $0 < \pi_k \leq 1$ and $\sum_{k=1}^g \pi_k = 1$, and $\alpha = (\alpha_1, \dots, \alpha_g)$ groups the parameters of the multinomial distributions with $\alpha_k = (\alpha_{k1}, \dots, \alpha_{kd})$ and $\alpha_{kj} = (\alpha_{kj1}, \dots, \alpha_{kj m_j})$, α_{kjh} being the probability that $\mathbf{x}^{jh} = 1$ conditionally to the component k .

We also define the mapping τ_{kj} from $\{1, \dots, m_j\}$ to $\{1, \dots, m_j\}$ which orders the modalities of \mathbf{x}^j by decreasing values of the probabilities α_{kjh} . For instance, $\tau_{kj}(1)$ gives the modality of \mathbf{x}^j having the largest probability α_{kjh} . By using the shorter notation $\alpha_{kj(h)} = \alpha_{kj \tau_{kj}(h)}$, we have $\alpha_{kj(h)} \geq \alpha_{kj(h+1)}$ ($1 \leq h < m_j$). Furthermore, since the multinomial distribution of \mathbf{x}^j has ℓ_{kj} modes then α_{kj} is defined in the constrained simplex $S(\ell_{kj}, m_j)$ where

$$S(\ell_{kj}, m_j) = \left\{ \alpha_{kj} : \sum_{h=1}^{m_j} \alpha_{kjh} = 1, \alpha_{kj(\ell_{kj}+1)} = \dots = \alpha_{kj(m_j)} \right\}. \quad (3)$$

In other words, uniformity holds for non-mode modalities. By using the other shorter notation $\mathbf{x}^{j(h)k} = \mathbf{x}^{j \tau_{kj}(h)}$, the conditional independence assumption between blocks involves the following pdf for the component k

$$p(\mathbf{x}; \alpha_k, \sigma, \ell_k) = \prod_{j=1}^d \prod_{h=1}^{m_j} (\alpha_{kjh})^{\mathbf{x}^{jh}} = \prod_{j=1}^d \prod_{h=1}^{m_j} (\alpha_{kj(h)})^{\mathbf{x}^{j(h)k}}. \quad (4)$$

Thus, CIM is included in CCM, since the conditional independence assumption between the initial variables is defined by putting $d = B$ and $\ell_{kj} = m_j - 1$. Indeed, in such case, each variable built a block, so $\sigma = (\{1\}, \dots, \{B\})$.

2.3. Conditional modes model characteristics

CMM is meaningful with its two levels of interpretation. Firstly, the intra-class dependencies of variables (equal between classes) are brought out by the repartition of the variables into blocks given by σ . Secondly, the intra-class and intra-block dependencies of modalities (possibly different between classes) are summarized by the modes (locations and probabilities). A shorter summary for each distribution is also available by using the following compact terms, defined on $[0, 1]$ and respectively reflecting the *complexity* and the *strength* of the intra-class and intra-block dependencies

$$\kappa_{kj} = \frac{\ell_{kj}}{m_j - 1} \text{ and } \rho_{kj} = \sum_{h=1}^{\ell_{kj}} \alpha_{kj(h)}. \quad (5)$$

For instance, the smaller is κ_{kj} and the larger is ρ_{kj} , the more massed in few characteristic modality crossings is the distribution, since the modes are interpreted as an over-contribution at the uniform distribution among all the modality crossings.

Note that the repartition of the variables into conditionally independent blocks identical between classes assures the model generic identifiability. Indeed, with this constraint, the results of [25] can be applied to prove the generic identifiability of the CMM (details are given in Appendix A). Despite the constraint of the same repartition of the variables into blocks for all the classes, the model stays flexible because of the specific block distribution.

The main idea of the former sparse versions of CIM (classical conditional independence model) proposed by [8] is to estimate only one mode for each multinomial distribution of the initial variable. Different constraints of equality are then added between the variables and/or classes. In fact, many of these models are included in the model family of CMM by putting $d = B$ and $\ell_{kj} = 1$. In addition, as CMM needs $\nu = (g - 1) + \sum_{k=1}^g \sum_{j=1}^d \ell_{kj}$ parameters, models of CMM's family can need less parameters than CIM—having $(g - 1) + g \times \sum_{b=1}^B (m_b - 1)$ parameters—although it takes into account the conditional dependencies.

2.4. New parametrization of the block distribution

The parsimonious versions of CIM introduced by [8] are meaningful since each multinomial is written with two parameters: one discrete giving the location of the mode of the distribution and one continuous giving its probability.

By using the same idea, we propose a new parametrization of the block distribution, denoted by $(\boldsymbol{\delta}_{kj}, \mathbf{a}_{kj})$, which facilitates the interpretation and the writing of the prior and posterior distributions of the block parameters (see Section 4). The discrete parameter $\boldsymbol{\delta}_{kj} = \{\delta_{kjh}; h = 1, \dots, \ell_{kj}\}$ determines the mode locations, since δ_{kjh} indicates the modality crossing where the mode h is located, with $\delta_{kjh} \neq \delta_{kjh'}$ if $h \neq h'$ and $\delta_{kjh} \in \{1, \dots, m_j\}$. The continuous parameter $\mathbf{a}_{kj} = (a_{kjh}; h = 1, \dots, \ell_{kj} + 1)$ determines the probability mass of the ℓ_{kj} modes by its first ℓ_{kj} elements (a_{kjh} with $h = 1, \dots, \ell_{kj}$) and the probability mass of the non-mode by its last element ($a_{kj\ell_{kj}+1}$). This parameter is defined on the truncated simplex involving $a_{kjh} \geq \frac{a_{kj\ell_{kj}+1}}{m_j - \ell_{kj}}$ ($h \leq \ell_{kj}$). The parameter $\boldsymbol{\alpha}_{kj}$ and the couple $(\boldsymbol{\delta}_{kj}, \mathbf{a}_{kj})$ are linked by

$$\alpha_{kjh} = \begin{cases} a_{kjh'} & \text{if } \exists h' \text{ such that } \delta_{kjh'} = h \\ \frac{a_{kj\ell_{kj}+1}}{m_j - \ell_{kj}} & \text{otherwise.} \end{cases} \quad (6)$$

3. Maximum likelihood estimate

The whole data set consisting of n independent and identically distributed individuals is denoted by $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_n)$. Remark that \mathbf{X} denotes the whole observed sample and not a random variable. The observed-data log-likelihood of CMM is also defined as

$$L(\boldsymbol{\theta}; \mathbf{X}, \boldsymbol{\omega}) = \sum_{i=1}^n \ln p(\mathbf{x}_i; \boldsymbol{\theta}, \boldsymbol{\omega}). \quad (7)$$

Since we use CMM to cluster, the indicator vector of the g classes denoted by $\mathbf{Z} = (\mathbf{z}_i; i = 1, \dots, n)$ with $\mathbf{z}_i = (z_{i1}, \dots, z_{ig})$ where $z_{ik} = 1$ if the individual \mathbf{x}_i arises from the class k and $z_{ik} = 0$ otherwise, is considered as a missing variable. The complete-data log-likelihood of CMM is then defined by

$$L(\boldsymbol{\theta}; \mathbf{X}, \mathbf{Z}, \boldsymbol{\omega}) = \sum_{i=1}^n \sum_{k=1}^g z_{ik} \ln (\pi_k p(\mathbf{x}_i; \boldsymbol{\alpha}_k, \boldsymbol{\sigma}, \ell_k)). \quad (8)$$

For the mixture models, the direct optimization on $\boldsymbol{\theta}$ to obtain the maximum likelihood estimate (MLE), denoted by $\hat{\boldsymbol{\theta}}$, involves solving equations having no analytical solution. So, we perform the parameter's estimation via an EM algorithm [26, 27], which is often simple and efficient for the missing data situation. It is an iterative algorithm which alternates between two

steps: the computation of the complete-data log-likelihood conditional expectation (E step) and its maximization (M step). At the iteration $[r]$, this algorithm is written as:

E step: conditional probabilities computation

$$t_{ik}(\boldsymbol{\theta}^{[r]}) = \frac{\pi_k^{[r]} p(\mathbf{x}_i; \boldsymbol{\alpha}_k^{[r]}, \boldsymbol{\sigma}, \ell_k)}{\sum_{k'=1}^g \pi_{k'}^{[r]} p(\mathbf{x}_i; \boldsymbol{\alpha}_{k'}^{[r]}, \boldsymbol{\sigma}, \ell_{k'})}. \quad (9)$$

M step: maximization of the complete-data log-likelihood

$$\pi_k^{[r+1]} = \frac{n_k^{[r]}}{n} \text{ and } \alpha_{kj(h)}^{[r+1]} = \begin{cases} \frac{n_{kj(h)}^{[r]}}{n_k^{[r]}} & \text{if } (1 \leq h \leq \ell_{kj}) \\ \frac{1 - \sum_{h'=1}^{\ell_{kj}} \alpha_{kj(h')}^{[r+1]}}{m_j - \ell_{kj}} & \text{otherwise,} \end{cases} \quad (10)$$

by using the notations $n_k^{[r]} = \sum_{i=1}^n t_{ik}(\boldsymbol{\theta}^{[r]})$ and $n_{kjh}^{[r+1]} = \sum_{i=1}^n t_{ik}(\boldsymbol{\theta}^{[r]}) x_i^{jh}$. Note that, at the M step of iteration $[r]$, the function τ_{kj} is redefined as the decreasing ordering function of the $n_{kjh}^{[r+1]}$ and allows to define $n_{kj(h)}^{[r+1]}$ with $n_{kj(h)}^{[r+1]} \geq n_{kj(h+1)}^{[r+1]}$.

4. Model selection via Metropolis-within-Gibbs sampler

The aim is to obtain the model $\hat{\omega}$ having the largest posterior probability. We assume that $p(g) = \frac{1}{g_{\max}}$ for $g = 1, \dots, g_{\max}$ and that $p(\boldsymbol{\sigma})$ (remind that g and $\boldsymbol{\sigma}$ are independent) and $p(\boldsymbol{\ell}|g, \boldsymbol{\sigma})$ follow uniform distributions. Let the g_{\max} models denoted by $\boldsymbol{\omega}^{(g)} = (g, \boldsymbol{\sigma}^{(g)}, \boldsymbol{\ell}^{(g)})$, for $g = 1, \dots, g_{\max}$, where $(\boldsymbol{\sigma}^{(g)}, \boldsymbol{\ell}^{(g)}) = \underset{\boldsymbol{\sigma}, \boldsymbol{\ell}}{\operatorname{argmax}} p(\boldsymbol{\sigma}, \boldsymbol{\ell} | \mathbf{X}, g)$. The best model is $\underset{g}{\operatorname{argmax}} p(\boldsymbol{\omega}^{(g)} | \mathbf{X})$ and it is found by applying the BIC approximation among those g_{\max} selected models. However, an exhaustive search strategy is not feasible for two correlated reasons: firstly, the number of couples $(\boldsymbol{\sigma}, \boldsymbol{\ell})$ can be excessively huge, and, secondly, the estimation of the MLE for each of them is an unnecessary waste of time computing. A Metropolis-within-Gibbs sampler strategy overcomes these two drawbacks at the same time, as we now describe.

For a fix value of g , the couple $(\boldsymbol{\sigma}^{(g)}, \boldsymbol{\ell}^{(g)})$ is estimated by the following Metropolis-within-Gibbs sampler [28] having $p(\boldsymbol{\sigma}, \boldsymbol{\ell} | g, \mathbf{X})$ as stationary

distribution and whose the iteration $[s]$ is written as

$$\boldsymbol{\theta}^{[s+1]} \sim p(\boldsymbol{\theta}|\boldsymbol{\omega}^{[s]}, \mathbf{X}, \mathbf{Z}^{[s]}) \quad (11)$$

$$\mathbf{Z}^{[s+1]} \sim p(\mathbf{Z}|\boldsymbol{\omega}^{[s]}, \mathbf{X}, \boldsymbol{\theta}^{[s+1]}) \quad (12)$$

$$(\boldsymbol{\sigma}^{[s+1]}, \boldsymbol{\ell}^{[s+1]}) \sim p(\boldsymbol{\sigma}, \boldsymbol{\ell}|\boldsymbol{\omega}^{[s]}, \mathbf{X}, \mathbf{Z}^{[s+1]}), \quad (13)$$

where $\boldsymbol{\omega}^{[s]} = (g, \boldsymbol{\sigma}^{[s]}, \boldsymbol{\ell}^{[s]})$. As the observed data are independent, the full conditional distribution of \mathbf{Z} is classical and is written as

$$p(\mathbf{Z}|\boldsymbol{\omega}, \mathbf{X}, \boldsymbol{\theta}) = \prod_{i=1}^n p(\mathbf{z}_i|\boldsymbol{\omega}, \mathbf{x}_i, \boldsymbol{\theta}) \text{ with } p(\mathbf{z}_i|\boldsymbol{\omega}, \mathbf{x}_i, \boldsymbol{\theta}) = \prod_{k=1}^g (t_{ik}(\boldsymbol{\theta}))^{z_{ik}}. \quad (14)$$

In this section, we firstly detail the full conditional distributions sampling the parameters (denoted by *instrumental elements*) by using block parametrization given in Section 2.4, and we secondly detail the sampling of $(\boldsymbol{\sigma}, \boldsymbol{\ell})$ (considered as the *interest elements*).

4.1. Sampling of the instrumental elements

We now detail the sampling according to $p(\boldsymbol{\theta}|\boldsymbol{\omega}^{[s]}, \mathbf{X}, \mathbf{Z}^{[s]})$ defined in (11).

Prior assumption. We assume the independence *a priori* between the class proportions and the block distribution parameters, involving that the prior of the whole parameter is written as

$$p(\boldsymbol{\theta}|\boldsymbol{\omega}) = p(\boldsymbol{\pi}|\boldsymbol{\omega}) \prod_{k=1}^g \prod_{j=1}^d p(\boldsymbol{\alpha}_{kj}|\boldsymbol{\omega}). \quad (15)$$

Note that this property of conditional independence is preserved for the distribution of $\boldsymbol{\theta}$ conditionally on $(\boldsymbol{\omega}, \mathbf{X}, \mathbf{Z})$, thus

$$p(\boldsymbol{\theta}|\boldsymbol{\omega}, \mathbf{X}, \mathbf{Z}) = p(\boldsymbol{\pi}|\boldsymbol{\omega}, \mathbf{X}, \mathbf{Z}) \prod_{k=1}^g \prod_{j=1}^d p(\boldsymbol{\alpha}_{kj}|\boldsymbol{\omega}, \mathbf{X}, \mathbf{Z}). \quad (16)$$

Prior and posterior distributions of $\boldsymbol{\pi}$. The Jeffreys non informative prior distribution, for a multinomial, is a conjugate Dirichlet distribution [29]. So, the prior and the posterior distributions of $\boldsymbol{\pi}$ [9] are respectively defined by

$$\boldsymbol{\pi}|\boldsymbol{\omega} \sim \mathcal{D}_g\left(\frac{1}{2}, \dots, \frac{1}{2}\right) \text{ and } \boldsymbol{\pi}|\boldsymbol{\omega}, \mathbf{X}, \mathbf{Z} \sim \mathcal{D}_g\left(\frac{1}{2} + n_1, \dots, \frac{1}{2} + n_g\right), \quad (17)$$

where $n_k = \sum_{i=1}^n z_{ik}$ (not equal to $n_k^{[r]}$).

Prior distribution of α_{kj} . We now use the parametrization of the block distribution $(\delta_{kj}, \mathbf{a}_{kj})$ (defined in Section 2.4). We assume the independence between the prior of δ_{kj} and of \mathbf{a}_{kj} , so

$$p(\alpha_{kj}|\omega) = p(\delta_{kj}|\omega)p(\mathbf{a}_{kj}|\omega). \quad (18)$$

We use a uniform distribution among all the mode locations and a conjugate truncated Dirichlet distribution³ as prior of \mathbf{a}_{kj} , so

$$p(\delta_{kj}|\omega) = \left(\frac{m_j}{\ell_{kj}}\right)^{-1} \text{ and } \mathbf{a}_{kj}|\omega \sim D_{\ell_{kj}+1}^t\left(\gamma_{kj1}, \dots, \gamma_{kj\ell_{kj}+1}; m_j\right), \quad (19)$$

where the γ_{kjh} are the parameters of the truncated Dirichlet distribution. In Appendix B, we justify why we now fix $\gamma_{kjh} = 1$. The proposed prior is also weakly informative since it is an uniform distribution.

Posterior distribution of α_{kj} . The posterior distribution of α_{kj} is written as

$$p(\alpha_{kj}|\omega, \mathbf{X}, \mathbf{Z}) = p(\delta_{kj}|\omega, \mathbf{X}, \mathbf{Z})p(\mathbf{a}_{kj}|\omega, \delta_{kj}, \mathbf{X}, \mathbf{Z}). \quad (20)$$

The distribution of $\delta_{kj}|\omega, \mathbf{X}, \mathbf{Z}$ is a multinomial one with too many values to be computable. Let the set $\tilde{\delta}_{kj} = \{\tilde{\delta}_{kjh}; h = 1, \dots, \ell_{kj}\}$ containing the indices of the ℓ_{kj} largest values of $n_{kjh} = \sum_{i=1}^n z_{ik} x_i^{jh}$ ordered such as

$$\forall h \in \{1, \dots, \ell_{kj} - 1\}, \quad n_{kj\tilde{\delta}_{kjh}} \geq n_{kj\tilde{\delta}_{kjh+1}}. \quad (21)$$

We assume that the difference between the mode probabilities and the non-mode probabilities are significant. So, we can approximate the full conditional distribution of δ_{kj} by a Dirac in $\tilde{\delta}_{kj}$. This approximation is strengthened by the fast convergence speed of the discrete parameters [30]. Concerning now \mathbf{a}_{kj} , as its prior is conjugated, its conditional distribution is explicitly defined as

$$\mathbf{a}_{kj}|\omega, \delta_{kj}, \mathbf{X}, \mathbf{Z} \sim \mathcal{D}_{\ell_{kj}+1}^t\left(1 + n_{kj(1)}, \dots, 1 + n_{kj(\ell_{kj})}, 1 + \bar{n}_{kj}^{\ell_{kj}}; m_j\right), \quad (22)$$

where $n_{kj(h)}$ is the h th larger value of the set $\{n_{kjh}; h = 1, \dots, m_j\}$ and $\bar{n}_{kj}^{\ell_{kj}} = n_k - \sum_{h=1}^{\ell_{kj}} n_{kj(h)}$.

³ $p(\mathbf{a}_{kj}|\omega) \propto \prod_{h=1}^{\ell_{kj}+1} (a_{kjh})^{\gamma_{kjh}-1} \mathbb{1}_{\left\{a_{kjh} \geq \frac{a_{kj\ell_{kj}+1}}{m_j - \ell_{kj}}\right\}}$.

4.2. Sampling of a new model according to $p(\boldsymbol{\sigma}, \boldsymbol{\ell} | \boldsymbol{\omega}^{[s]}, \mathbf{X}, \mathbf{Z}^{[s+1]})$

The sampling of $\boldsymbol{\omega}^{[s+1]} = (g, \boldsymbol{\sigma}^{[s+1]}, \boldsymbol{\ell}^{[s+1]})$ according to Equation (13) is performed in two steps. Firstly, a new repartition of the variables into blocks and the mode number of the modified blocks, respectively denoted by $\boldsymbol{\sigma}^{[s+1]}$ and $\boldsymbol{\ell}^{[s+1/2]}$, are sampled by one iteration of a Metropolis-Hastings algorithm. Secondly, the mode number of each block is sampled by one MCMC iteration. Thus, the sampling of $\boldsymbol{\omega}^{[s+1]}$ is decomposed into the two following steps

$$(\boldsymbol{\sigma}^{[s+1]}, \boldsymbol{\ell}^{[s+1/2]}) \sim p(\boldsymbol{\sigma}, \boldsymbol{\ell} | \boldsymbol{\omega}^{[s]}, \mathbf{X}, \mathbf{Z}^{[s+1]}) \quad (23)$$

$$\boldsymbol{\ell}^{[s+1]} \sim p(\boldsymbol{\ell} | \boldsymbol{\omega}^{[s+1/2]}, \mathbf{X}, \mathbf{Z}^{[s+1]}), \quad (24)$$

where $\boldsymbol{\omega}^{[s+1/2]} = (g, \boldsymbol{\sigma}^{[s+1]}, \boldsymbol{\ell}^{[s+1/2]})$. Thus, this chain has $p(\boldsymbol{\sigma}, \boldsymbol{\ell} | g, \mathbf{X}, \mathbf{Z}^{[s+1]})$ as stationary distribution.

4.2.1. Metropolis-Hastings algorithm to sample $\boldsymbol{\omega}^{[s+1/2]}$

The sampling of $\boldsymbol{\omega}^{[s+1/2]}$ is performed by one iteration of the Metropolis-Hastings algorithm divided into two steps. Firstly, the proposal distribution $q(\cdot; \boldsymbol{\omega}^{[s]})$ generates a candidate $\boldsymbol{\omega}^* = (g, \boldsymbol{\sigma}^*, \boldsymbol{\ell}^*)$. Secondly $\boldsymbol{\omega}^{[s+1]}$ is sampled according to the acceptance probability $\mu^{[s]}$ defined by

$$\mu^{[s]} = 1 \wedge \frac{p(\mathbf{X}, \mathbf{Z}^{[s]} | \boldsymbol{\omega}^*)}{p(\mathbf{X}, \mathbf{Z}^{[s]} | \boldsymbol{\omega}^{[s]})} \frac{q(\boldsymbol{\omega}^{[s]}; \boldsymbol{\omega}^*)}{q(\boldsymbol{\omega}^*; \boldsymbol{\omega}^{[s]})}. \quad (25)$$

The computation of $\mu^{[s]}$ involves to compute the integrated complete-data likelihood. In Section 4.2.3, we described how to solve this problem without using the biased BIC approximation or using too much time computing MCMC methods. The sampling of $\boldsymbol{\omega}^{[s+1/2]}$ is written as

$$\begin{aligned} \boldsymbol{\omega}^* &\sim q(\cdot; \boldsymbol{\omega}^{[s]}) \\ \boldsymbol{\omega}^{[s+1/2]} &= \begin{cases} \boldsymbol{\omega}^* & \text{with a probability } \mu^{[s]} \\ \boldsymbol{\omega}^{[s]} & \text{otherwise.} \end{cases} \end{aligned} \quad (26)$$

The proposal distribution $q(\cdot; \boldsymbol{\omega}^{[s]})$ samples $\boldsymbol{\omega}^*$ in two steps. The first step changes the block affectation of one variable. In practice, $\boldsymbol{\sigma}^*$ is uniformly sampled in $V(\boldsymbol{\sigma}^{[s]}) = \{\boldsymbol{\sigma} : \exists! b \text{ as } b \in \boldsymbol{\sigma}_j^{[s]} \text{ and } b \notin \boldsymbol{\sigma}_j\}$. The second step uniformly samples the mode numbers among all its possible values for the modified blocks while $\ell_{kj}^* = \ell_{kj}^{[s]}$ for non-modified blocks (i.e. j as $\boldsymbol{\sigma}_j^{[s]} = \boldsymbol{\sigma}_j^*$).

4.2.2. MCMC algorithm to sample $\ell^{[s+1]}$

This step allows to increase or decrease the mode number of each block by one at each iteration. So, $\ell_{kj}^{[s+1]}$ is sampled according to $p(\ell_{kj}|\boldsymbol{\omega}^{[s+1/2]}, \mathbf{X}, \mathbf{Z}^{[s]})$ defined by

$$p(\ell_{kj}|\boldsymbol{\omega}^{[s+1/2]}, \mathbf{X}, \mathbf{Z}^{[s+1]}) \propto \begin{cases} p(\mathbf{X}^j|\mathbf{Z}^{[s+1]}, \ell_{kj}) & \text{if } |\ell_{kj} - \ell_{kj}^{[s+1/2]}| < 2 \\ & \text{and } \ell_{kj} \notin \{0, m_j\}. \\ 0 & \text{otherwise.} \end{cases} \quad (27)$$

Thus, this algorithm needs the value of $p(\mathbf{X}^j|\mathbf{Z}, \ell_{kj})$ defined by

$$p(\mathbf{X}^j|\mathbf{Z}, \ell_{kj}) = \int_{S(\ell_{kj}, m_j)} \prod_{h=1}^{m_j} (\alpha_{kjh})^{n_{kjh}} d\boldsymbol{\alpha}_{kj}. \quad (28)$$

That we have to detail now.

4.2.3. The integrated complete-data likelihood

The integrated complete-data likelihood is defined as

$$p(\mathbf{X}, \mathbf{Z}|\boldsymbol{\omega}) = p(\mathbf{Z}|\boldsymbol{\omega}) \prod_{k=1}^g \prod_{j=1}^d p(\mathbf{X}^j|\mathbf{Z}, \ell_{kj}), \quad (29)$$

where $\mathbf{X}^j = (\mathbf{X}_i^j; i = 1, \dots, n)$. Note that the quantities $p(\mathbf{X}, \mathbf{Z}|\boldsymbol{\omega})$ and $p(\mathbf{X}^j|\mathbf{Z}, \ell_{kj})$ are respectively needed to compute the acceptance probability of the Metropolis-Hastings algorithm (see Equation (25)) and to sample the number of modes (see Equation (27)) and can be evaluated by BIC-like approximations [31, 32]. For instance, the integrated complete-data likelihood is approximated by

$$\ln p(\mathbf{X}, \mathbf{Z}|\boldsymbol{\omega}) = \ln p(\mathbf{X}, \mathbf{Z}|\boldsymbol{\theta}^*, \boldsymbol{\omega}) - \frac{\nu}{2} \ln n + \mathcal{O}(1), \quad (30)$$

$\boldsymbol{\theta}^*$ begin the maximum complete-data likelihood estimate. However, this kind of approximation is only asymptotically true and over-estimates the mode numbers (see Section 5.1). As $\mathbf{Z}|\boldsymbol{\omega}$ follows a uniform distribution among all the possible partitions, we propose to compute each $p(\mathbf{X}^j|\mathbf{Z}, \ell_{kj})$ to obtain $p(\mathbf{X}, \mathbf{Z}|\boldsymbol{\omega})$. This computation is not easy since $\boldsymbol{\alpha}_{kj}$ is defined on $S(\ell_{kj}; m_j)$ and not on the whole simplex of size ℓ_{kj} (except when $\ell_{kj} = m_j - 1$, in such case we can use the approach of CIM [9]). An explicit formula is given in the

following proposition whose the proof is given in Appendix B by performing an exact computation of the integral over the continuous parameters and an approximation on the discrete ones.

Proposition 1. *The integrated complete-data likelihood is approximated, by neglecting the sum over the discrete parameters of the modes locations and by performing the exact computation on the continuous parameters, by*

$$p(\mathbf{X}^j | \mathbf{Z}, \ell_{kj}) \approx \left(\frac{1}{m_j - \ell_{kj}} \right)^{\bar{n}_{kj}^{\ell_{kj}}} \prod_{h=1}^{\ell_{kj}} \frac{Bi\left(\frac{1}{m_j - h + 1}; n_{kj(h)} + 1; \bar{n}_{kj}^h + 1\right)}{m_j - h}, \quad (31)$$

where $Bi(x; a, b) = B(1; a, b) - B(x; a, b)$, $B(x; a, b)$ being the incomplete beta function defined by $B(x; a, b) = \int_0^x w^a (1 - w)^b dw$.

From the previous expression, it is straightforward to obtain $p(\mathbf{X}, \mathbf{Z} | \boldsymbol{\omega})$.

5. Simulations

5.1. Integrated complete-data likelihood: comparison of both approaches

Aim. During this experiment, we highlight the biases of the BIC criterion for the selection of the number of modes and the gain provided by the proposed computation of the integrated complete-data likelihood.

Data generation. As we want to compare both approaches for the selection of the number of modes, we simulate samples composed by n i.i.d individuals arisen from a multinomial distribution per modes $\mathcal{M}_s(r, r, r, \frac{1-3r}{s-3}, \dots, \frac{1-3r}{s-3})$ with s modalities and three modes having a probability r . For different sizes of sample, 10^5 samples are generated with different values of (r, s) .

Results. Figure 1 gives a comparison between the proposed approach and the BIC-like approximation for the selection of the number of modes. The proposed criterion obtains best results than the BIC criterion in the four studied situations for the large size of sample. Furthermore, it allows to never overestimates the mode number. Finally, its variability is smaller than the BIC criterion one. We enter now into more specific comments.

In case (a), modes have a large probability mass and they are easily detected since there are few modalities. Thus, both criteria have the same behavior since they find the true number of modes with a probability close to one even for small samples.

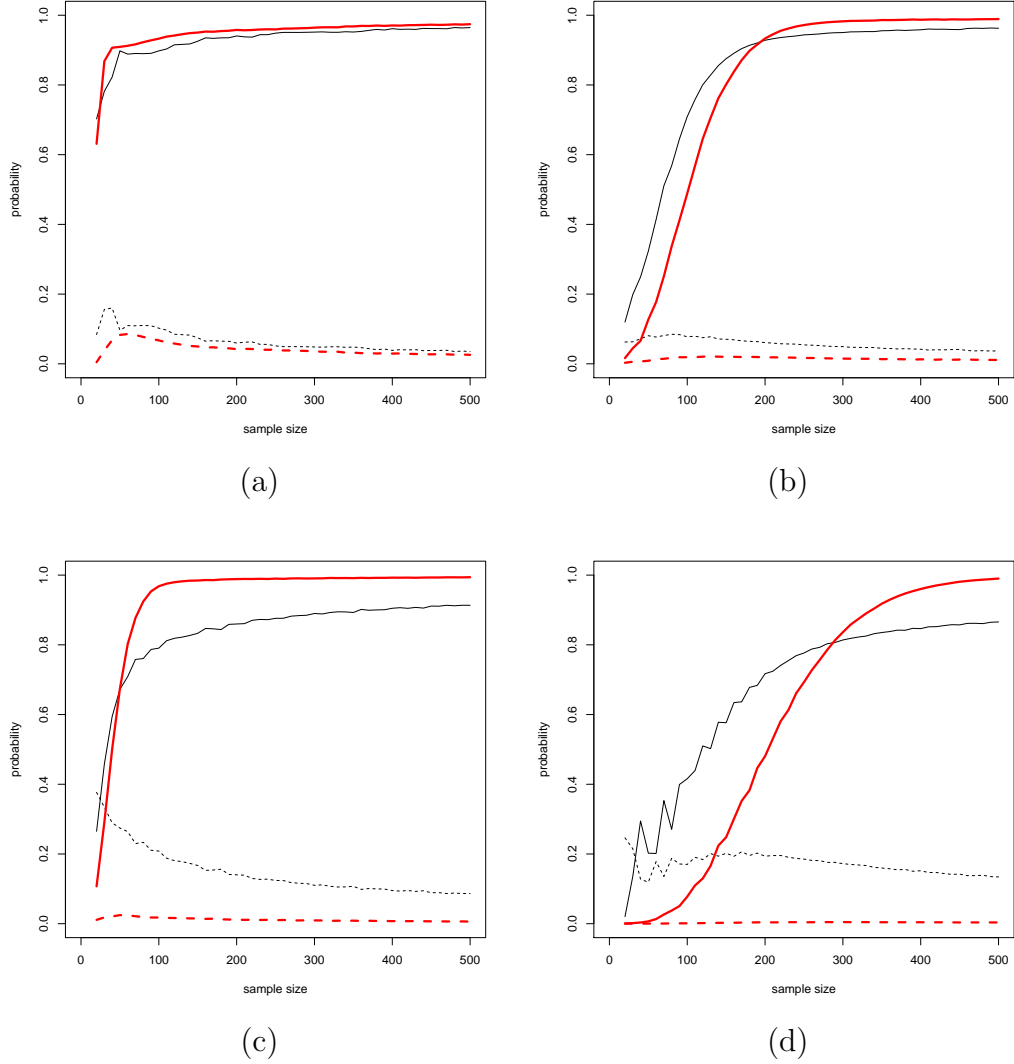


Figure 1: Probability that the BIC criterion (represented in fine black lines) and the proposed approach (represented in bold red lines) select the true number of modes (represented in plain line) and over-estimate (represented in dotted line) (a) $r=0.3$, $s=9$; (b) $r=0.2$, $s=9$; (c) $r=0.2$, $s=18$; (d) $r=0.1$, $s=27$.

When the mode probabilities decrease (case (b)), it is more difficult to identify them. In such case, the BIC criterion allows to better find the true number of modes with a moderate overestimation risk, for the small samples

(size lower than 150), than the proposed approach which can underestimates the number of modes. When the sample size is larger than 200, the proposed approach obtains better results since it finds a true number of modes almost always while the BIC criterion keeps an overestimation risk.

If the number of modalities increases (case (c)), then the problem becomes harder and the proposed approach also shows its interest since the BIC criterion is strongly biased in such case. The BIC criterion keeps this bias even for a large data set while the proposed approach almost always finds the true number of modes when the sample size is larger than 100.

Finally, note that in the more complex situations like in case (d) (few probability mass for the modes and large number of modalities), the proposed approach underestimates the number of modes when the sample size is small then converges to the true mode values when the sample size increases. Note that, in such case, the bias of the BIC criterion keeps significant even for a large data set.

Based on this experiment, the proposed criterion seems most relevant since its asymptotic behavior is better than the BIC criterion, it never overestimates the mode number and its variability the smaller than the BIC criterion.

5.2. Simulation with well specified model

Aim. During this experiment, we highlight the good behavior of the algorithms (Metropolis-within-Gibbs sampler and EM algorithm) for performing the model selection and the estimation of the MLE. Thus, data are generated according to CMM, then the model and the MLE are estimated. The quality of the estimation is determined by the Kullback-Leibler divergence. We show that this divergence converges to zero when the sample size increases. So, we conclude to the good behavior of both algorithms.

Data generation. A data set of six variables with three modalities is generated according to a bi-component CMM with the following parameters: $\sigma = (\{1, 2\}, \{3, 4\}, \{5, 6\})$, $\ell_{kj} = 2$, $\pi = (0.5, 0.5)$, $\alpha_{kj} = (0.4, 0.4, 0.2/7)$, where the modes are located at different modality crossings for both classes.

Results. For different values of $n = (50, 100, 200, 400, 800)$, 100 samples are generated. The Kullback-Leibler divergence is computed between the true and the estimated parameters. Table 1 presents the mean of this divergence.

n	50	100	200	400	800
mean	0.656	0.117	0.061	0.028	0.015
sd	0.636	0.052	0.018	0.007	0.003

Table 1: Mean and standard deviation of the Kullback-Leibler divergence computed between the true parameters of the specified model and the maximum likelihood estimates associated to the model selected by the Metropolis-within-Gibbs algorithm for different sample size.

As the Kullback-Leibler divergence converges to zero, when the sample size increases, we claim that the estimated distribution converges to the true one. Thus, we conclude to the good behavior of the estimation algorithm.

5.3. Simulation with misspecified model

Aim. During this experiment, we underline that the flexibility of CMM allows it to keep good results even if the model is misspecified. Thus, we simulate samples according to a bi-component mixture model where the intra-class dependencies are different for both components. A tuning parameter allows us to modify the strength of the intra-class dependencies and the class overlapping. The results of CMM are compared to those of CIM.

Data generation. A data set of size 100 is sampled from the following bi-component mixture model of dimension six

$$p(\mathbf{x}; \boldsymbol{\theta}) = 0.5 \prod_{h=1}^3 p(\mathbf{x}^{2h-1}, \mathbf{x}^{2h}; \boldsymbol{\theta}) + 0.5 p(\mathbf{x}^1; \boldsymbol{\theta}) p(\mathbf{x}^6; \boldsymbol{\theta}) \prod_{h=1}^2 p(\mathbf{x}^{2h}, \mathbf{x}^{2h+1}; \boldsymbol{\theta}), \quad (32)$$

with $p(\mathbf{x}^j, \mathbf{x}^{j+1}; \boldsymbol{\theta}) = p(\mathbf{x}^j; \boldsymbol{\theta}) (\lambda \mathbb{1}_{\{\mathbf{x}^j = \mathbf{x}^{j+1}\}} + (1 - \lambda) p(\mathbf{x}^{j+1}; \boldsymbol{\theta}))$ and with $\mathbf{x}^j \sim \mathcal{M}_3(1/3, 1/3, 1/3)$. Thus, when $\lambda = 0$, the sample is generated by a uniform distribution and classes are confused. The larger is the tuning parameter λ , the larger are the intra-class dependencies and the class separation. Note that CMM is not the good model since the conditionally correlated variables are not the same in both classes.

Results. For different values of $\lambda = (0.2, 0.4, 0.6, 0.8)$, 100 samples are generated. The Kullback-Leibler divergence associated to the model with the best number of classes (selected by the BIC criterion among $g = 1, 2, 3, 4$) is

λ	0.2	0.4	0.6	0.8
CMM	0.09 (1.00)	0.25 (1.16)	0.53 (2.08)	0.87 (2.10)
CIM	0.11 (1.00)	0.27 (1.00)	1.67 (1.12)	5.79 (1.40)

Table 2: Kullback-Leibler divergence and mean of the class number obtained by CMM and CIM.

computed. Table 2 presents the results obtained by CMM and CIM.

The larger is λ , the larger is the Kullback-Leibler divergence for both models. However, the flexibility of CMM allows to keep an acceptable value of the Kullback-Leibler divergence while this divergence grows dramatically faster with CIM. Furthermore, when the classes are well separated (large value of λ), CMM finds more often the true class number than CIM.

6. Applications

For both applications, the estimation of CMM was performed by the R package `CoModes`. Both data set are available in `CoModes` developed by the authors. Appendix C displays the R code of the second application and can be used as a tutorial of `CoModes`.

6.1. Seabirds clustering

Data. We study a biological data set describing 153 puffins (seabirds) by five plumage and external morphological characteristics presented in Table 3 [33]. These seabirds are divided into three subspecies *dichrous* (84 birds), *lherminieri* (34 birds) and *subalaris* (35 birds).

variables	m_j	modalities				
collar	5	none	continuous
eyebrows	4	none	very pronounced
sub-caudal	4	white	black	black and white	BLACK and white	
border	3	none	...	many		
gender	2	male	female			

Table 3: Presentation of the five plumage and external morphological variables describing the puffins.

Experimental settings. The subspecies memberships of the individuals are blinded. For $g = 1, \dots, 6$, the MLE of CIM is obtained by 25 initializations of an EM algorithm while 25 chains of 3000 iterations are performed for the model selection of CMM followed by 25 initializations of EM algorithm to find the MLE.

Results. Table 4 presents the values of the BIC criterion for both models and different class numbers. Even if both models select two components, the values of the BIC criterion are better for CMM than for CIM for all the number of classes. Thus, CMM better fits the data than CIM.

g	1	2	3	4	5	6
CMM	-711	-691	-701	-709	-721	-727
CIM	-711	-706	-722	-745	-775	-805

Table 4: Values of the BIC criterion for different class numbers and for CMM and CIM. Boldface indicates the best values of this criterion.

According to Table 5 displaying the confusion matrix between the estimated partitions and the subspecies, we claim that the Subalaris are more different than the two other subspecies. Indeed, both models affect all the Subalaris in class 2. If the estimated partitions by both models are similar, we remark that CMM affects less other subspecies in this class than CIM.

	CMM		CIM	
	class 1	class 2	class 1	class 2
Dichrous	52	32	48	36
Lherminieri	23	11	22	12
Subalaris	0	35	0	35

Table 5: Confusion tables between the subspecies and estimated partition into two classes.

Figure 2(a) displays the seabirds on the first correspondence analysis plan and indicates the subspecies. We note that all the Subalaris are in the same location (bottom left) for the first principal correspondence map. We display the partition corresponding to the best model (CMM with two components) in Figure 2(b). Note that, for both model, the first principal correspondence axe allows to define a classification rule.

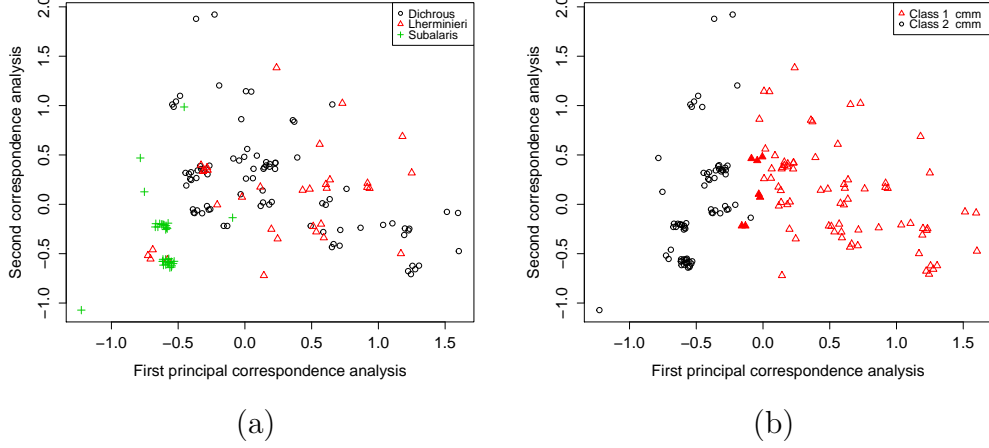


Figure 2: Seabirds on the first principal correspondence analysis map (a) with the sub-species and (b) with the best CMM estimated partition. The bold triangles indicate the individuals affected in class 1 for CMM and in class 2 for CIM. An i.i.d. uniform noise on $[0, 0.1]$ has been added on both axes for each individual in order to improve visualization.

We now describe the best bi-component model of CMM. Even if the estimated model assumes the conditional independence between the variables, this model is of interest because of its sparsity. Indeed, it is more parsimonious than CIM since a small number of modes is estimated as shown by the summary proposed by κ_{kj} and ρ_{kj} defined in (5) and presented in Table 6. Thus, the first variables are characterized by few modalities with a high probability. As the variables are conditionally independent, the κ_{kj} indicates the number of modalities having a probability upper than the uniform distribution. For example the multinomial distribution of the variable sub-caudal has two modes for both classes (so $\kappa_{kj} = 2/3$).

	collar	eyebrows	sub-caudal	border	gender
class 1	0.75 (0.93)	0.67 (0.91)	0.67 (0.88)	1.00 (1.00)	1.00 (0.55)
class 2	0.75 (0.98)	0.67 (0.77)	0.67 (0.99)	0.50 (0.97)	1.00 (0.57)

Table 6: Summary of the CMM with three classes: κ_{kj} is displayed in plain and ρ_{kj} is displayed in parenthesis.

The maximum likelihood estimates of the component parameters are dis-

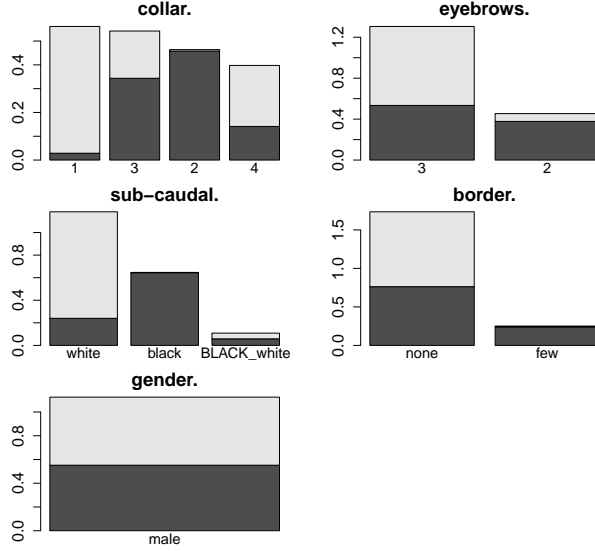


Figure 3: Class parameters of the bi-components CMM estimated on the Seabirds data. The black color (respectively the gray color) corresponds to the probability mass of the modes for the class 1 (respectively to the class 2).

played in Figure 3. Each sub-figure corresponds to a block of variable, thus we note again that the estimated model assumes the conditional independence. For each block of variables, the modality crossings where one mode is estimated for at least one component are focused. For these modality crossings, we display their cumulated probability masses for each component (the component are identifiable by different colors). These modality crossings are presented by decreasing order of cumulated probability mass. Note that the mode locations are discriminative since the modality black (resp. white) has a probability of 0.64 (resp. 0.24) for the class 1 while the modality white (resp. BLACK and white) has a probability of 0.94 (resp. 0.05).

Finally, the conditional independence assumption seems realistic since the conditional Cramer's V measures, presented in Table 7, are small. We also perform a bootstrap test of the global nullity of the Cramer's V by generating 1000 samples. We obtain a p-value of 0.91, so the conditional independence assumption is validated.

1	0.14	0.15	0.23	0.21	1	0.14	0.09	0.11	0.28
	1	0.36	0.20	0.13		1	0.24	0.21	0.26
		1	0.13	0.19			1	0.02	0.07
			1	0.01				1	0.17
				1					1
(a) Class 1					(b) Class 2				

Table 7: Matrix of the Cramer’s V measures computed according to the estimated classes.

6.2. Acute inflammations clustering

A tutorial of the R package **CoModes** performing the clustering of the Acute inflammations data set is presented in Appendix C.

Data. We want to cluster 120 patients [34] described by five binary variables (occurrence of nausea (Nau), lumbar pain (Lum), urine pushing (Pus), micturition pains (Mic) and burning of urethra (Bur)) and by one three modalities variables (temperature of the patient (Tem): $T < 37C$, $37C \leq T < 38C$ and $38C \leq T$). We know that some patients have one of the following diseases of the urinary system: inflammation of urinary bladder and Nephritis of renal pelvis origin.

Experimental conditions. We use the same experimental conditions as the Seabirds clustering.

Results. Table 8 presents the values of the BIC criterion for both models and different class numbers. For each class number, the BIC criterion value of CMM is better than for CIM. Furthermore, CMM selects three classes while CIM selects four classes. This phenomenon can be due to the violated conditional independence assumption of CIM.

g	1	2	3	4	5	6
CMM	-510	-351	-338	-345	-399	-401
CIM	-527	-478	-439	-407	-412	-418

Table 8: Values of the BIC criterion for different classes number and for CMM and CIM. Boldface indicates the best values of this criterion.

Note that the estimated distributions of CIM and CMM are different. The obtained partition are also different. Table 9 displays the confusion matrices

between the best model of CMM and the models of CIM with three and four classes. Thus, if 29 individuals constitute a group which is well separated of the other individuals (class 3) for the three models, the other individuals have a class membership determined by the selected model.

	CMM				CMM		
	c1	c2	c3		c1	c2	c3
CIM c1	40	0	0	CIM c1	40	0	0
CIM c2	10	41	0	CIM c2	10	20	0
CIM c3	0	0	29	CIM c3	0	21	0
				CIM c4	0	0	29

Table 9: Confusion matrices between the best model of CMM and the models of CIM with three and four classes.

Figure 4 displays the individuals on the 1-5 principal correspondence analysis map where the estimated classes are well separated.




Figure 4: Individuals on the 1-5 principal correspondence analysis map with the best CMM estimated partition. An i.i.d. uniform noise on $[0, 0.1]$ has been added on both axes for each individual in order to improve visualization. Colors and symbols indicate the class membership.

The model CMM with three classes has the following repartition of the variables into blocks: $\sigma = (\{\text{Tmp}, \text{Pus}, \text{Mic}, \text{Bur}\}, \{\text{Nau}\}, \{\text{Lum}\})$. As shown

by the summary ρ_{kj} and κ_{kj} displayed in Table 10, the three classes are concentrated in few modality crossings for the block one and in one location with a probability close to one for the two other blocks.

	Tmp, Nau, Lum, Mic	Pus	Bur
Class 1	0.41 (1.00)	1.00 (1.00)	1.00 (0.99)
Class 2	0.33 (0.99)	1.00 (1.00)	1.00 (1.00)
Class 3	0.25 (0.99)	1.00 (1.00)	1.00 (1.00)

Table 10: Summary of the CMM with three classes: κ_{kj} is displayed in plain and ρ_{kj} is displayed in parenthesis.

The following class interpretation is based on the class parameters displayed by Figure 5. Note that the variables *urine pushing* and *burning of urethra* are the most discriminative ones.

- The majority class (42%) groups individuals having no nausea and no lumber pain.
- The second class (34%) groups individuals having no nausea but lumber pain.
- The third class (24%) groups individuals having nausea and lumber pain. Furthermore, these individuals have some fever and micturition pain.

7. Conclusion

In this article, we have presented a new mixture model (CMM) to cluster categorical data. Its strength is to relax the conditional independence assumption and to stay parsimonious. A summary of the distribution is given by κ_{kj} and ρ_{kj} while each class can be summarized by the mode locations. As shown by the Seabirds application, CMM can improve the results of the classical latent class model even if the conditional independence assumption is true, thank's to its sparsity.

The combinatorial problems of the block detection and of the modes number selection is solved by a Metropolis-within-Gibbs algorithm and use the computation of the integrated complete-data likelihood. Thus, this approach can be used to select the interactions of the log-linear mixture model




Figure 5: Estimated parameters of the tri-component CMM displayed by the barplot function of the package **CoModes**. Black color corresponds to class 1, black gray color corresponds to class 2 and pale gray color corresponds to class 3.

per block. The R package **CoModes** allows to perform the model selection and the parameter's estimation. Both data sets presented in this article are included in this package. To efficiently reduce the computing time, the functions of this package will be soon implemented in C++.

However, the model is hardly estimated if the data set has a large number of variables. Some constraints on the block variables repartition could also be added (for instance the number of variables into blocks could be limit at three variables). Another solution could be to estimate the model by a forward/backward strategy but it is know that these method are sub-optimal.

Finally, we imposed the equality of the repartition of the variables into blocks for all the classes. This property allows us to prove the generic identifiability of CMM. This lack of flexibility is counterbalanced by flexible block distribution. However, one could try to relax the class-equality of σ with the model no-identifiability risk.

References

- [1] K. Jajuga, A. Sokołowski, H. Bock, Classification, clustering and data analysis: recent advances and applications, Springer Verlag, 2002.
- [2] J. Huang, M. Ng, H. Rong, Z. Li, Automated variable weighting in k-means type clustering, Pattern Analysis and Machine Intelligence, IEEE Transactions on 27 (5) (2005) 657–668.
- [3] M. Chavent, V. Kuentz, J. Saracco, A partitioning method for the clustering of categorical variables, in: Classification as a Tool for Research, Springer, 2010, pp. 91–99.
- [4] C. Guinot, J. Latreille, D. Malvy, P. Preziosi, P. Galan, S. Hercberg, M. Tenenhaus, Use of multiple correspondence analysis and cluster analysis to study dietary behaviour: food consumption questionnaire in the SU. VI. MAX. cohort, European journal of epidemiology 17 (6) (2001) 505–516.
- [5] G. Govaert, Data analysis, Vol. 136, Wiley. com, 2010.
- [6] G. McLachlan, D. Peel, Finite mixutre models, Wiley Series in Probability and Statistics: Applied Probability and Statistics, Wiley-Interscience, New York, 2000.
- [7] L. Goodman, Exploratory latent structure analysis using both identifiable and unidentifiable models, Biometrika 61 (2) (1974) 215–231.
- [8] G. Celeux, G. Govaert, Clustering criteria for discrete data and latent class models, Journal of classification 8 (2) (1991) 157–176.
- [9] C. Biernacki, G. Celeux, G. Govaert, Exact and Monte Carlo calculations of integrated likelihoods for the latent class model, Journal of Statistical Planning and Inference 140 (11) (2010) 2991–3002.
- [10] D. Hand, K. Yu, Idiot’s Bayes Not So Stupid after All?, International Statistical Review 69 (3) (2001) 385–398.
- [11] B. Reboussin, E. Song, A. Shrestha, K. Lohman, M. Wolfson, A latent class analysis of underage problem drinking: Evidence from a community sample of 16–20 year olds, Drug and alcohol dependence 83 (3) (2006) 199–209.

- [12] S. Strauss, D. Rindskopf, J. Astone-Twerell, D. Des Jarlais, H. Hagan, Using latent class analysis to identify patterns of hepatitis C service provision in drug-free treatment programs in the US, *Drug and alcohol dependence* 83 (1) (2006) 15–24.
- [13] P. Van Hattum, H. Hoijsink, Market Segmentation Using Brand Strategy Research: Bayesian Inference with Respect to Mixtures of Log-Linear Models, *Journal of Classification* 26 (3) (2009) 297–328.
- [14] J. Vermunt, Multilevel latent class models, *Sociological methodology* 33 (1) (2003) 213–239.
- [15] J. Vermunt, Multilevel mixture item response theory models: an application in education testing, *Proceedings of the 56th session of the International Statistical Institute. Lisbon, Portugal* (2007) 22–28.
- [16] Y. Qu, M. Tan, M. Kutner, Random Effects Models in Latent Class Analysis for Evaluating Accuracy of Diagnostic Tests, *Biometrics* 52 (3) (1996) pp. 797–810.
- [17] I. Gollini, T. Murphy, Mixture of latent trait analyzers for model-based clustering of categorical data, *Statistics and Computing* (2013) 1–20.
- [18] D. Bartholomew, M. Knott, I. Moustaki, *Latent variable models and factor analysis: A unified approach*, Vol. 899, Wiley. com, 2011.
- [19] A. Agresti, *Categorical data analysis*, Vol. 359, John Wiley and Sons, 2002.
- [20] H. Bock, *Loglinear models and entropy clustering methods for qualitative data, Classification as a tool of research*. North Holland, Amsterdam (1986) 19–26.
- [21] J. Hagenaars, Latent structure models with direct effects between indicators local dependence models, *Sociological Methods & Research* 16 (3) (1988) 379–405.
- [22] M. Espeland, S. Handelman, Using Latent Class Models to Characterize and Assess Relative Error in Discrete Measurements, *Biometrics* 45 (2) (1989) pp. 587–599.

- [23] M. Jorgensen, L. Hunt, Mixture model clustering of data sets with categorical and continuous variables, in: Proceedings of the Conference ISIS, Vol. 96, 1996, pp. 375–384.
- [24] L. Hunt, M. Jorgensen, Theory & Methods: Mixture model clustering using the MULTIMIX program, Australian & New Zealand Journal of Statistics 41 (2) (1999) 154–171.
- [25] E. Allman, C. Matias, J. Rhodes, Identifiability of parameters in latent structure models with many observed variables, The Annals of Statistics 37 (6A) (2009) 3099–3132.
- [26] A. Dempster, N. Laird, D. Rubin, Maximum likelihood from incomplete data via the EM algorithm, Journal of the Royal Statistical Society. Series B (Methodological) (1977) 1–38.
- [27] G. McLachlan, T. Krishnan, The EM algorithm, Wiley Series in Probability and Statistics: Applied Probability and Statistics, Wiley-Interscience, New York, 1997.
- [28] C. Robert, G. Casella, Monte Carlo statistical methods, Springer Verlag, 2004.
- [29] C. Robert, The Bayesian choice: from decision-theoretic foundations to computational implementation, Springer, 2007.
- [30] C. Choirat, R. Seri, Estimation in discrete parameter models, Statistical Science 27 (2) (2012) 278–293.
- [31] G. Schwarz, Estimating the dimension of a model, Annals of Statistics 6 (1978) 461–464.
- [32] E. Lebarbier, T. Mary-Huard, Une introduction au critre BIC : fondements thoriques et interprétation, Journal de la SFdS 147 (1) (2006) 39–57.
- [33] V. Bretagnolle, Personal communication, source: Museum.
- [34] J. Czerniak, H. Zarzycki, Application of rough sets in the presumptive diagnosis of urinary system diseases, Artificial Intelligence and Security in Computing Systems, ACS’2002 9th International Conference Proceedings (2003) 41–51.

- [35] J. Kruskal, Three-way arrays: rank and uniqueness of trilinear decompositions, with application to arithmetic complexity and statistics, *Linear algebra and its applications* 18 (2) (1977) 95–138.
- [36] J. Kruskal, More factors than subjects, tests and treatments: An indeterminacy theorem for canonical decomposition and individual differences scaling, *Psychometrika* 41 (3) (1976) 281–293.

Appendix A. Generic identifiability of CMM

We demonstrate that CMM is generically identifiable, *i.e.* the parameter space where the model is not identifiable has a Lebesgue measure equal to zero. To do this, we adapt the demonstration of the generic identifiability of CMM given by [25] and based on the Kruskal theorem [35, 36]. By using the conditional independence between the blocks of variables, we present a sufficient condition for the generic identifiability of CMM which is a relation between the class number and the mode number. As this demonstration is just an adaptation of the proof given by [25], some technical details are not reminded here. The demonstration is cut into three steps: we start by a reminder of the Kruskal results for the three-way tables, then we show the generic identifiability of CMM with three blocks of variables and we finish by an extension to CMM with more than three blocks.

Kruskal results. For a matrix M , the Kruskal rank of M , denoted by $\text{rank}_K M$ is the largest number I such that every set of I rows of M are linearly independent.

THEOREM 1 (Kruskal [35, 36]). Let $I_j = \text{rank}_K M_j$. If

$$I_1 + I_2 + I_3 \geq 2g + 2,$$

then the tensor $[M_1, M_2, M_3]$ uniquely determines the M_j , up to simultaneous permutation and rescaling rows.

Generic identifiability of CMM with three blocks. Let $k_0 = \underset{k}{\operatorname{argmin}} \ell_{kj}$ and the matrix M_j where

$$M_j(k, h) = \alpha_{kj\tau_{k_0j}(h)}. \quad (\text{A.1})$$

By denoting by $\xi_j = \min_k \ell_{kj} + 1$, generically, we have

$$\text{rank}_K M_j = \min(g, \xi_j).$$

COROLLARY 1 The parameters of CMM with three blocs are generically identifiable, up to label swapping, provided:

$$\min(g, \xi_1) + \min(g, \xi_2) + \min(g, \xi_3) \geq 2g + 2.$$

Generic identifiability of CMM with more than three blocks. In the same way that [25], we generalize the result with d blocks by observing that d blocks of categorical variables can be combined into three categorical variables. Thus, we can apply the Kruskal theorem.

COROLLARY 2 We consider a CMM with d blocks where $d \leq 3$. If there exists a tri-partition of the set $\{1, \dots, d\}$ into three disjoint non empty subsets S_1 , S_2 and S_3 , such that $\gamma_i = \prod_{j \in S_i} \xi_j$ with

$$\min(g, \gamma_1) + \min(g, \gamma_2) + \min(g, \gamma_3) \geq 2g + 2, \quad (\text{A.2})$$

then the model parameters are generically identifiable up to label swapping.

Appendix B. Proof of Proposition 1

In this Section, a proof of Proposition 1 is given. We firstly define a new parametrization of the block distribution facilitating the integrate complete-data likelihood computation. We secondly define the prior distribution of the new block parametrization according to the other parametrization. Thirdly, we underline the relation between the embedded models. We conclude by the integrate complete-date likelihood computation, which is the target result.

Appendix B.1. New parametrization of the block distribution

Without loss of generality, we assume that the elements of δ_{kj} are ordered by decreasing values of the probability mass associated to them and we introduce the new parametrization of \mathbf{a}_{kj} denoted $\boldsymbol{\varepsilon}_{kj}$ where $\boldsymbol{\varepsilon}_{kj} \in \mathcal{E}_{kj} = \left[\frac{1}{m_j}; 1\right] \times \dots \times \left[\frac{1}{m_j - \ell_{kj}}; 1\right]$ and where ε_{kjh} is defined by

$$\varepsilon_{kjh} = \begin{cases} a_{kj} \delta_{kjh} & \text{if } h = 1 \\ \frac{a_{kj} \delta_{kjh}}{\prod_{h'=1}^{h-1} (1 - \varepsilon_{kjh'})} & \text{otherwise.} \end{cases}$$

Lemma 1. *The conditional probability of \mathbf{X}^j is*

$$p(\mathbf{X}^j | \mathbf{Z}, \ell_{kj}, \tilde{\boldsymbol{\delta}}_{kj}, \boldsymbol{\varepsilon}_{kj}) = \prod_{h=1}^{\ell_{kj}} (\varepsilon_{kjh})^{n_{kj(h)}} (1 - \varepsilon_{kjh})^{\bar{n}_{kj}^h}, \quad (\text{B.1})$$

Proof.

$$\begin{aligned}
p(\mathbf{X}^j | \mathbf{Z}, \ell_{kj}, \tilde{\boldsymbol{\delta}}_{kj}, \boldsymbol{\varepsilon}_{kj}) &= p(\mathbf{X}^j | \mathbf{Z}, \ell_{kj}, \boldsymbol{\alpha}_{kj}) \\
&= \prod_{h=1}^{m_j} (\alpha_{kjh})^{n_{kjh}} \\
&= \left[\prod_{h=1}^{\ell_{kj}} (\alpha_{kj(h)})^{n_{kj(h)}} \right] \alpha_{kj(\ell_{kj}+1)}^{\bar{n}_{kj}^{\ell_{kj}}} \\
&= \varepsilon_{kj1}^{n_{kj(1)}} \prod_{h=2}^{\ell_{kj}} \left[\varepsilon_{kjh}^{n_{kj(h)}} \left(\prod_{h'=1}^{h-1} (1 - \varepsilon_{kjh'})^{n_{kj(h')}} \right) \right] \prod_{h=1}^{\ell_{kj}} (1 - \varepsilon_{kjh})^{\bar{n}_{kj}^{\ell_{kj}}} \\
&= \prod_{h=1}^{\ell_{kj}} (\varepsilon_{kjh})^{n_{kj(h)}} (1 - \varepsilon_{kjh})^{\bar{n}_{kj}^h}.
\end{aligned}$$

□

Appendix B.2. Prior distribution

Lemma 2. *The prior distribution of $\boldsymbol{\varepsilon}_{kj}$ is*

$$p(\boldsymbol{\varepsilon}_{kj} | \boldsymbol{\omega}, \boldsymbol{\delta}_{kj}) = \frac{m_j}{m_j - \ell_{kj}}. \quad (\text{B.2})$$

Proof. We remind that $\mathbf{a}_{kj} | \boldsymbol{\omega} \sim D_{\ell_{kj}+1}^t(1, \dots, 1; m_j)$ and that

$$p(\mathbf{a}_{kj}, \boldsymbol{\delta}_{kj} | \boldsymbol{\omega}) = p(\boldsymbol{\alpha} | \boldsymbol{\omega}) = p(\boldsymbol{\varepsilon}_{kj}, \boldsymbol{\delta}_{kj} | \boldsymbol{\omega}). \quad (\text{B.3})$$

So, we deduce the pdf of the prior distribution of $\boldsymbol{\varepsilon}_{kj}$

$$p(\boldsymbol{\varepsilon}_{kj} | \boldsymbol{\delta}_{kj}, \boldsymbol{\omega}) = \frac{\prod_{h=1}^{\ell_{kj}} (\varepsilon_{kjh})^{\gamma_{kjh}-1} (1 - \varepsilon_{kjh})^{\sum_{h'=h+1}^{\ell_{kj}+1} (\gamma_{kjh'}-1)}}{\int_{\boldsymbol{\varepsilon}_{kj} \in \mathcal{E}_{kj}} \prod_{h=1}^{\ell_{kj}} (\varepsilon_{kjh})^{\gamma_{kjh}-1} (1 - \varepsilon_{kjh})^{\sum_{h'=h+1}^{\ell_{kj}+1} (\gamma_{kjh'}-1)} d\boldsymbol{\varepsilon}_{kj}}. \quad (\text{B.4})$$

Thus, each ε_{kjh} follows a truncated Beta distribution on the parameters space $\left[\frac{1}{m_j - h + 1}, 1 \right]$ denoted by $\mathcal{Be}(\gamma_{kjh}, \sum_{h'=h+1}^{\ell_{kj}+1} (\gamma_{kjh'} - 1) + 1)$. To assure the positivity of the parameters of the truncated Beta distributions, we put $\gamma_{kjh} = 1$, so

$$p(\boldsymbol{\varepsilon}_{kj} | \boldsymbol{\delta}_{kj}, \boldsymbol{\omega}) = \frac{m_j}{m_j - \ell_{kj}}. \quad (\text{B.5})$$

□

Appendix B.3. Relation between embedded models

Lemma 3. Let the model with ℓ_{kj}^\ominus modes and the parameters $(\tilde{\boldsymbol{\delta}}_{kj}^\ominus, \boldsymbol{\varepsilon}_{kj}^\ominus)$ and let the model with ℓ_{kj} modes and the parameters $(\tilde{\boldsymbol{\delta}}_{kj}, \boldsymbol{\varepsilon}_{kj})$. Both modes are defined as such that $\ell_{kj}^\ominus = \ell_{kj} - 1$, that the ℓ_{kj}^\ominus modes having the largest probabilities have the same locations ($\forall h \in \boldsymbol{\delta}_{kj}^\ominus, h \in \boldsymbol{\delta}_{kj}$) and the same probability masses ($\varepsilon_{kjh}^\ominus = \varepsilon_{kjh}, h < \ell_{kj}$). These embedded models follow this relation

$$\frac{p(\mathbf{X}^j | \mathbf{Z}, \ell_{kj}, \tilde{\boldsymbol{\delta}}_{kj}, \boldsymbol{\varepsilon}_{kj})}{p(\mathbf{X}^j | \mathbf{Z}, \ell_{kj}^\ominus, \tilde{\boldsymbol{\delta}}_{kj}^\ominus, \boldsymbol{\varepsilon}_{kj}^\ominus)} = \frac{(m_j - \ell_{kj} + 1)^{\tilde{n}_{kj}^{\ell_{kj}-1} - 1}}{(m_j - \ell_{kj})^{\tilde{n}_{kj}^{\ell_{kj}}}} (\varepsilon_{\ell_{kj}})^{n_{kj}(\ell_{kj})} (1 - \varepsilon_{\ell_{kj}})^{\tilde{n}_{kj}^{\ell_{kj}}}. \quad (\text{B.6})$$

Proof. We start by the following relation

$$\frac{p(\mathbf{X}^j | \mathbf{Z}, \ell_{kj}, \boldsymbol{\alpha}_{kj})}{p(\mathbf{X}^j | \mathbf{Z}, \ell_{kj}^\ominus, \boldsymbol{\alpha}_{kj}^\ominus)} = \frac{\alpha_{kj\ell_{kj}}^{n_{kj}(\ell_{kj})} (\alpha_{kj\ell_{kj}+1})^{\tilde{n}_{kj}^{\ell_{kj}}}}{\alpha_{kj\ell_{kj}}^{\ominus \tilde{n}_{kj}^{\ell_{kj}-1}}}. \quad (\text{B.7})$$

Note that, $\varepsilon_{kjh} = \varepsilon_{kjh}^\ominus$ when $(h = 1, \dots, \ell_{kj} - 1)$, since $\alpha_{kj(h)} = \alpha_{kj(h)}^\ominus$ and $\tilde{\tau}_{\ell_{kj}}(h) = \tilde{\tau}_{\ell_{kj}-1}(h)$ when $(h = 1, \dots, \ell_{kj} - 1)$. Then, by using the reparametrization in $\boldsymbol{\varepsilon}_{kj}$, the proof is completed. \square

Appendix B.4. Integrated complete-data likelihood

The integrated complete-data likelihood is finally approximated, by neglecting the sum over the discrete parameters of the modes locations and by performing the exact computation on the continuous parameters, by

$$p(\mathbf{X}^j | \mathbf{Z}, \ell_{kj}) \approx \left(\frac{1}{m_j - \ell_{kj}} \right)^{\tilde{n}_{kj}^{\ell_{kj}}} \prod_{h=1}^{\ell_{kj}} \frac{Bi\left(\frac{1}{m_j - h + 1}; n_{kj(h)} + 1; \tilde{n}_{kj}^h + 1\right)}{m_j - h}, \quad (\text{B.8})$$

where $Bi(x; a, b) = B(1; a, b) - B(x; a, b)$, $B(x; a, b)$ being the incomplete beta function defined by $B(x; a, b) = \int_0^x w^a (1 - w)^b dw$. From the previous expression, it is straightforward to obtain $p(\mathbf{X}, \mathbf{Z} | \boldsymbol{\omega})$.

Proof of Proposition 1. If, for the model with $\ell_{kj} - 1$ modes, the best modes locations are known and given by $\tilde{\boldsymbol{\delta}}_{kj}^\ominus$ then the conditional probability of \mathbf{X}^j

for a model with ℓ_{kj} modes is

$$p(\mathbf{X}^j | \mathbf{Z}, \ell_{kj}, \tilde{\boldsymbol{\delta}}_{kj}^{\ominus}, \boldsymbol{\varepsilon}_{kj}) = \frac{1}{m_j - \ell_{kj} + 1} \sum_{\tau \in \{1, \dots, m_j\} \setminus \{\tilde{\boldsymbol{\delta}}_{kj}^{\ominus}\}} p(\mathbf{X}^j | \mathbf{Z}, \ell_{kj}, \{\tilde{\boldsymbol{\delta}}_{kj}^{\ominus}, \tau\}, \boldsymbol{\alpha}_{kj}^{\ominus}, \boldsymbol{\varepsilon}_{kj}), \quad (\text{B.9})$$

Thus, by approximating this sum by its maximum element, we obtain that

$$p(\mathbf{X}^j | \mathbf{Z}, \ell_{kj}, \tilde{\boldsymbol{\delta}}_{kj}^{\ominus}, \boldsymbol{\varepsilon}_{kj}) \approx \frac{1}{m_j - \ell_{kj} + 1} p(\mathbf{X}^j | \mathbf{Z}, \ell_{kj}, \tilde{\boldsymbol{\delta}}_{kj}, \boldsymbol{\alpha}_{kj}^{\ominus}, \boldsymbol{\varepsilon}_{kj}). \quad (\text{B.10})$$

By using the proposition 3, we obtain that:

$$\frac{p(\mathbf{X}^j | \mathbf{Z}, \ell_{kj}, \tilde{\boldsymbol{\delta}}_{kj}^{\ominus}, \boldsymbol{\varepsilon}_{kj})}{p(\mathbf{X}^j | \mathbf{Z}, \ell_{kj}^{\ominus}, \tilde{\boldsymbol{\delta}}_{kj}^{\ominus}, \boldsymbol{\varepsilon}_{kj}^{\ominus})} \approx \frac{(m_j - \ell_{kj} + 1)^{\bar{n}_{kj}^{\ell_{kj}-1} - 1}}{(m_j - \ell_{kj})^{\bar{n}_{kj}^{\ell_{kj}}}} (\varepsilon_{\ell_{kj}})^{n_{kj}(\ell_{kj})} (1 - \varepsilon_{\ell_{kj}})^{\bar{n}_{kj}^{\ell_{kj}}}. \quad (\text{B.11})$$

As $p(\mathbf{X}^j | \mathbf{Z}, \ell_{kj} = 0) = (m_j)^{-n_k}$, by applying recursively the previous expression, we obtain that

$$p(\mathbf{X}^j | \mathbf{Z}, \ell_{kj}, \boldsymbol{\varepsilon}_{kj}) \approx \left(\frac{1}{m_j - \ell_{kj}} \right)^{\bar{n}_{kj}^{\ell_{kj}}} \prod_{h=1}^{\ell_{kj}} \frac{(\varepsilon_{kj h})^{n_{kj}(h)} (1 - \varepsilon_{kj h})^{\bar{n}_{kj}^h}}{m_j - h + 1}. \quad (\text{B.12})$$

□

Appendix C. Acute inflammation data set clustering with the R package CoModes

```
# Package loading
> require(CoModes)

# Loading of the data set Acute
> data(acute)

# Discretization of the first variable to obtain categorical variables
> acute[acute[,1] < 37, 1] <- 1
> acute[acute[,1] > 38, 1] <- 3
> acute[acute[,1] > 3, 1] <- 2
> acute[,1] <- factor(acute[,1], levels=1:3,
```

```

labels=c("cold","normal","fever"))

# Model selection and parameter estimation of CMM
>res.CoModes <- CoModescluster(acute[,1:6],2)

# Summary of the model
# Table like Table 10
>summary(res.CoModes)

# Plot of the parameters
# Barplot like Figure 5
>barplot(res.CoModes)

# Plot of the individuals in a multiple correspondence analysis map,
# the colors and symbols indicate the class membership estimated by CMM.
# Plot like Figure 4
>plot(res.CoModes,c(1,5))

```